

IN THE SPECIFICATION:

Please revise the following paragraphs to read as follows:

[0004] System and Methods for Intrusion Detection with Dynamic Window Sizes, application number 10/208,402, filed July 30, 2002, now U.S. Patent 7,162,741, which claims priority to application number 60/308,621, filed July 30, 2001.

[0011] System and Methods for Adaptive Model Generation for Detecting Intrusions in Computer Systems, application number 10/352,342, filed January 27, 2003, now U.S. Patent 7,225,343, which claims the benefit of application number 60/351,913, filed January 25, 2002.

[0017] Malicious activity may arise from any computer within the network or from multiple computers acting in concert. Typically, the malicious entity launches TCP/IP-based “probes,” which are attempts to connect with targeted network devices. “Scans” are systematic groups of probes originating from a single source computer or group of collaborating sources. Scans and probes are executed by malicious users and worms to find opportunities to attack or break ~~in to~~ into a targeted victim computer, and typically precede the actual attack. The attack itself may be an attempt to breach the security to the computer to obtain, e.g., user identification, access codes or other proprietary information or to interfere with the operation of the computer, e.g., by overloading resources, or redirecting the processing capabilities of the computer. All of these activities – probes, scans and attacks- are view as security threats. As will be appreciated, it is advantageous to detect probing and scanning sources to forewarn of a likely subsequent attempt to attack or break ~~in to~~ into a target victim computer.

[0019] Prior art intrusion detection methods include misuse detection and anomaly detection. Misuse detection requires *a priori* knowledge of an attack pattern. Online activity is evaluated with respect of a model of the malicious behavior, and activity that is consistent with the misuse model is flagged. Misuse detection offers the advantage of requiring relatively low computational resources. However, attack signatures must be known, and the misuse model must be designed to encompass all possible variations of the pertinent attack. Unfortunately, malicious users and programmers who write new worms often discover new ways to attack that are not known to programmers who write signature rules to detect attacks; and, as a result, IDSs often do not detect these attacks.

[0026] The prior art approaches are thus incapable of automatically and effectively detecting slow, stealthy surveillance activities or new, unknown threats, or significant variations of known threats. Accordingly, there is a need for new IDS methods and tools that can automatically detect, characterize and enable effective response to new threats without consuming inordinate computational and human resources.

[0027] In one aspect, the present invention is a generalized scan/probe and attack detector that overcomes the above limitations through the use of machine learning techniques, and by correlating detection alerts among multiple detectors. The generalized detector comprises a stealthy scan/probe detector, a standard IDS and an anomaly detector. Each alert includes IP address information about the source and target of the malicious activity, and this information is used to correlate the information from each of the detectors.

[0031] Once generated, the connection models are used by a detector to monitor online activity to detect malicious surveillance behavior without any requirement for a priori knowledge of system behavior. ~~The invention adapts to changes in the network and applications.~~

[0032] In another aspect, the present invention correlates scanning and probing activities with alerts generated by the anomaly detector, and may also correlate with the alerts generated by the intrusion detection system (IDS).

[0038] FIG. 2 is a schematic depiction of an illustrative embodiment of the automated detection and adaptive model learning methods of the present invention.

[0041] FIG. 5 is a Venn diagram useful in understanding the relationship between the scan/probe detector, intrusion detector and anomaly detector[.];

[0044] FIG. 8 is a schematic depiction of an illustrative embodiment of the computer program of FIG. 6; ~~and~~

[0045] FIG. 9 is a schematic illustration of a display screen used in an embodiment of the invention[.]; and

[0046] FIGs. 10-16 depict performance statistics related to an application of the invention in monitoring surveillance attacks on a large computer network.

[0051] In a preferred embodiment, automated scan/probe detection method 202 stores statistics on source IP addresses that initiate probes and scans over long periods of time. The Scan/Probe Detector maintains these statistics and IP addresses on a “watch list” of source IP scanners. After sufficient evidence (as determined by a user-selectable parameter selectable by a user) is gathered of a scanning activity, an “alert” is generated detailing this behavior. The alert is updated as new evidence is gathered. These alerts

are provided to an analyst. They may also be “correlated” with other alerts generated by other detectors. “Correlation” means that the respective alerts are aggregated combined if they contain a “common IP address,” either as a source or destination IP address.

[0052] The second main process of this aspect of the invention is known as adaptive learning 204 in Fig. 2. Adaptive learning makes use of knowledge from security audits performed by human analysts. Such audits generally involve analysis of correct positive detections, false positive detections, correct negative detections, and false alarms, negative detections. This process improves intrusion detection performance by automatically building improved detection models with, e.g., improved ability to accurately classify new intrusions[.,,]. Event classification models 216 and filtering operations 220 are refined according to behavior-based model generation and feature discovery process 240, based upon generated worthwhile alert notifications 222 and archives of the data used to generate them. The output of step 240 comprises candidate models 242, which are further refined through testing and quality assurance (QA) step 244. Step 244 produces interim refined models 248 for evaluation by analysts using a workstation 250. After refinement is complete, production models 246 are available for use by automated detection process 202. Testing / QA step 244 may optionally function with simulated raw data 143, processed by offline versions of sensing step 144 and event classification 216.

[0070] Address<1.1.1.1,2.2.2.2,TCP,111,222> would need to be matched with <2.2.2.2,1.1.1.1,TCP,222,111> in order to recognize that they represented the same connection. Instead the addresses are associated with same identifier for this connection no matter if address 1.1.1.1 sends the packet or address 2.2.2.2 sends the packet. Flags

fields in the connection extrapolator thereafter record from which direction interesting events (i.e., first packet, etc.) initiate.

[0071] Fig. 4 shows an embodiment of cost-based filtering of alerts to prioritize important alerts in order to simplify security analysts' tasks. Output from sensing step 144 are processed by alert prefilter 410 to remove redundancy. The prefiltered alerts 218 are further processed by correlation filter 430 according to cost model 420. This cost model prioritizes alerts according to such criteria as severity of the attack, importance of the network components or data affected, and the cost of preemptive action. The important alert stream 460 is delivered to the intrusion response team 228.

[0077] Fig. 8 shows a specific embodiment of computer program product ~~500~~ 600 of Fig. 6 5. Sensing step 144 supplies event information 214, which may be temporarily stored in buffer 710 and output on path 712. Worker 720 utilizes production models stored in memory 780 to perform initial event evaluation, yielding raw alerts 722. The raw alerts may be temporarily stored in buffer 730. Worker 740 processes raw alerts to detect probes and/or scans or to produce other alerts. Parameters for this worker may also be stored in memory 780. Filter 750 applies a cost-based filter to further process alerts 742. Output 752 may be stored in buffer 770, and subsequently processed by visualization station 224 or report generation station 790. Job submission/spool manager 510 orchestrates the operation of the aforementioned elements.

[0079] Fig. 9 depicts a display screen in an illustrative embodiment of a user interface used in the practice of the invention. The screen includes four displays: a display 910 of the activity of the top ten threats, a display 920 of the top targets as identified by all three probe/scan, IDS and anomaly detectors, a display 930 of further

details about selected threats and a display 940 of further details about selected targets. As indicated in Fig. 9, display 910 provides a plot of attack severity level with time. As shown in Fig. 9, the time period is three weeks but other time periods may be selected by the system user. Likewise, attack severity may be plotted and displayed for individual attackers or individual targets or for groups of attackers or groups of targets as selected by the system user. Display 910 has the advantage of making trends readily apparent at a glance at the display. Display 920 provides the IP address URL of the target and an estimate of the severity of the attack in the form of a numerical score. Advantageously, the score is the score computed by step 320. Displays 930 and 940 provide additional information in the form of an indication of the country, domain and source of the attacker and target, respectively. As depicted in Fig. 9, this display shows further details about the most severe attackers and most severely attacked targets, but these displays can be scrolled to provide more information about less severe events as well.